Netlist's Technology Tutorial

Netlist Inc.
v.
Micron Technology, Inc. et al

Case No. 2:22-CV-203-JRG-RSP

Patents-in-Suit

Distributed Data Buffer Patents

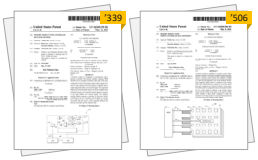

U.S. Patent 10,949,339

U.S. Patent 10,860,506

On-Module Power Management Integrated Circuit "PMIC" Patents

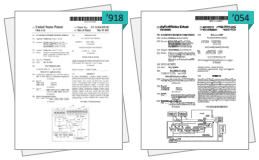

U.S. Patent 11,016,918

U.S. Patent 11,232,054

High Bandwidth Memory "HBM" Patents

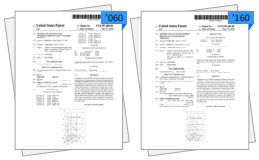

U.S. Patent 8,787,060

U.S. Patent 9,318,160

Technology Tutorial Topics

- **Overview**
- DDR4 LRDIMM
- DDR5's on-Module Power Management
- Vertically stacked high bandwidth memory

Components of a Computer Server

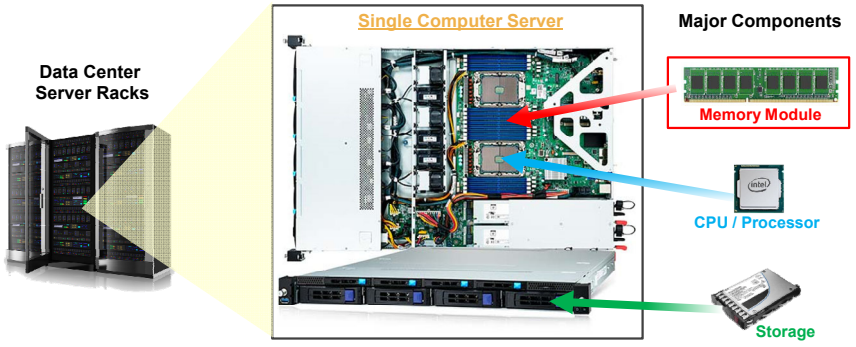

- CPU
- Storage
- Memory
- Input / Output

Memory Module Terminology

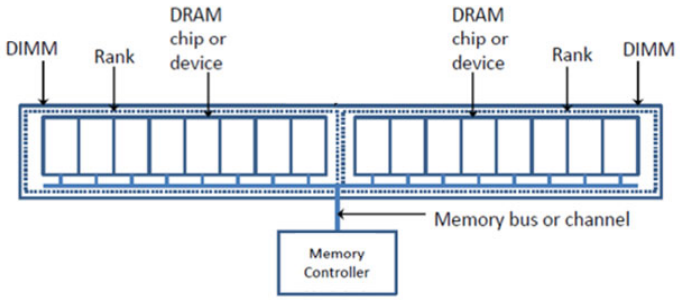

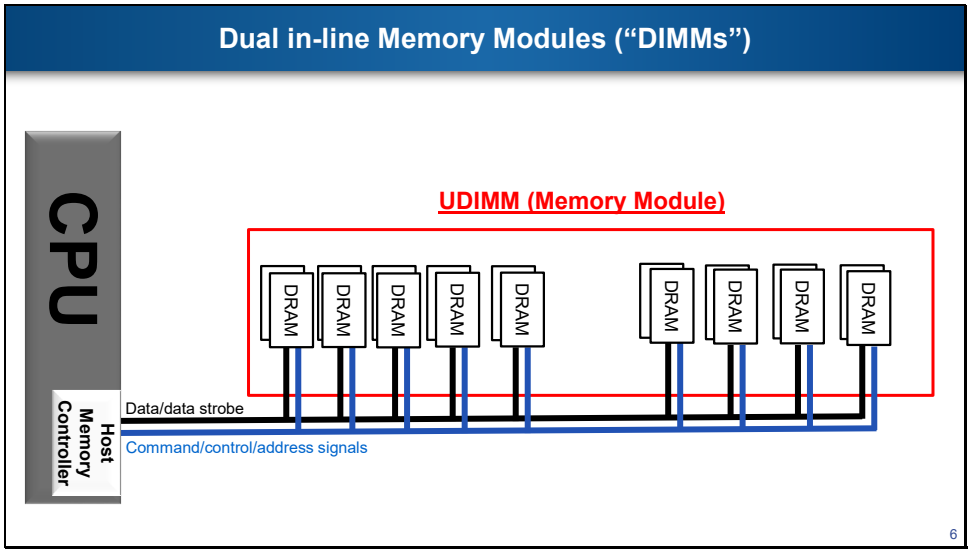

Dual in-line memory modules "DIMMs" are generally the most common type of memory module found in desktop computers and computer servers. A DIMM is a printed circuit board with memory devices (e.g., dynamic RAM or "DRAM") mounted on the printed circuit board along with other important components. The most simple form of DIMM is the UDIMM or unbuffered/unregistered DIMMs, illustrated above. In a UDIMM, a portion of the host CPU called the memory controller sends to and receives data directly from the DRAM components on the modules. The memory controller also sends command, control, address, and clock signals to the DRAM components.

See Johan De Gelas, *An Overview of Server DIMM types* (August 3, 2012), *available at* https://www.anandtech.com/show/6068/lrdimms-rdimms-supermicros-latest-twin/2.

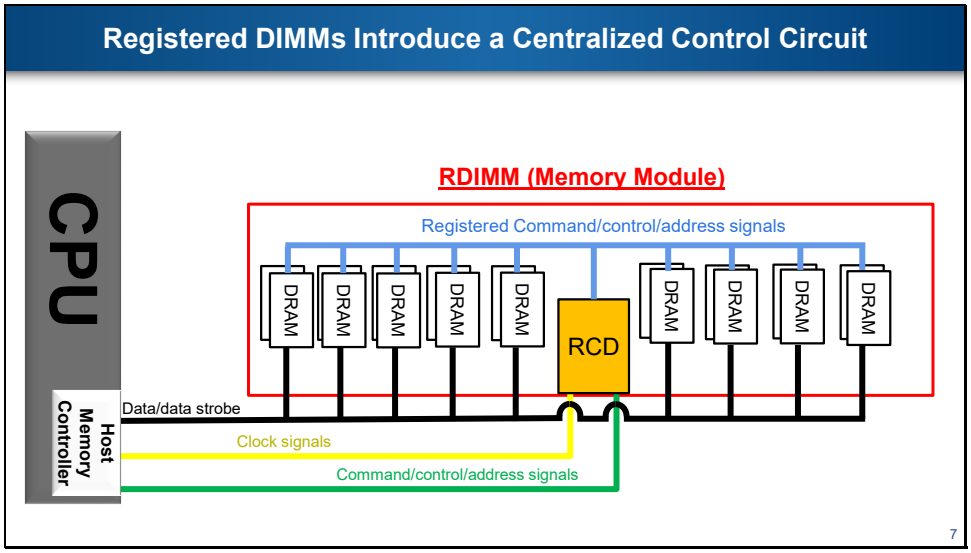

Another type of DIMM is a Registered DIMM (RDIMM) which features a control circuit on the DIMM that acts as a "mediator" between the memory controller and the DRAMs. This control circuit is called a Registering Clock Driver or "RCD." The RCD receives all the command, control, address, and clock signals sent by the host memory controller and then outputs the registered command, control, address, and clock signals to the DRAMs.

This differs from UDIMM, where the command, control, address, and clock signals are sent to each individual DRAM. In a UDIMM, the memory controller is in direct electrical communication with all the memory devices on the module, which exert a load on the memory controller; as the number of memory devices increase, so too does the load, negatively impacting system performance.

Because in an RDIMM, all the command, control, address, and clock signals are sent to a single control unit (the RCD) the host memory controller experiences, or "sees," only the loading of the RCD for command, control, address, and clock signals. This greatly reduces the apparent loading for command, control, address, and clock signals, which improves signal quality. In RDIMMs, data signals are still sent to each individual DRAM, however.

See Johan De Gelas, *An Overview of Server DIMM types* (August 3, 2012), *available at* https://www.anandtech.com/show/6068/lrdimms-rdimms-supermicros-latest-twin/2.

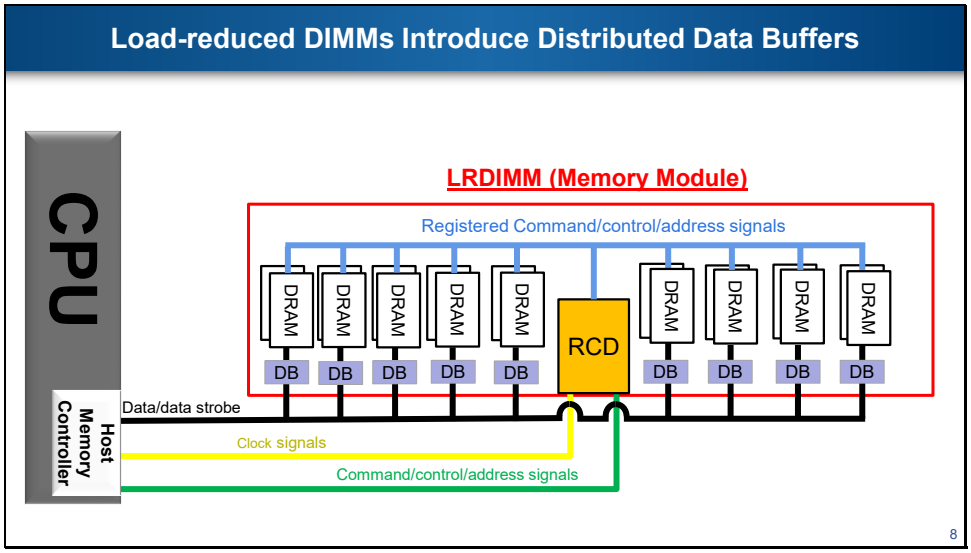

Another type of DIMM is a Load-reduced DIMM ("LRDIMM"), which features multiple data buffers (each abbreviated as a "DB") in addition to an RCD. Each data buffer is placed between a DRAM's data signal and data strobe signal lines on the one hand and the host memory controller on the other hand. Placing a data buffer between the signal lines for data and data strobe of the DRAMs and the host memory controller helps reduce the load experienced by the memory controller.

LRDIMMs incorporated distributed data buffers to reduce data loads experienced by the memory controller because generally, loading can be an issue for data, too, especially when there is more than one DRAM component on a data line. With the data buffers, the host memory controller experiences (or, "sees") only the loading of the data buffer on each data line. As a result, LRDIMMs not only reduce the load associated with the command, control, address and clock signals, but also the load associated with the data signals. This allows LRDIMMs to further improve the signal quality over a comparable RDIMMs.

See Johan De Gelas, *An Overview of Server DIMM types* (August 3, 2012), *available at* https://www.anandtech.com/show/6068/lrdimms-rdimms-supermicros-latest-twin/2.

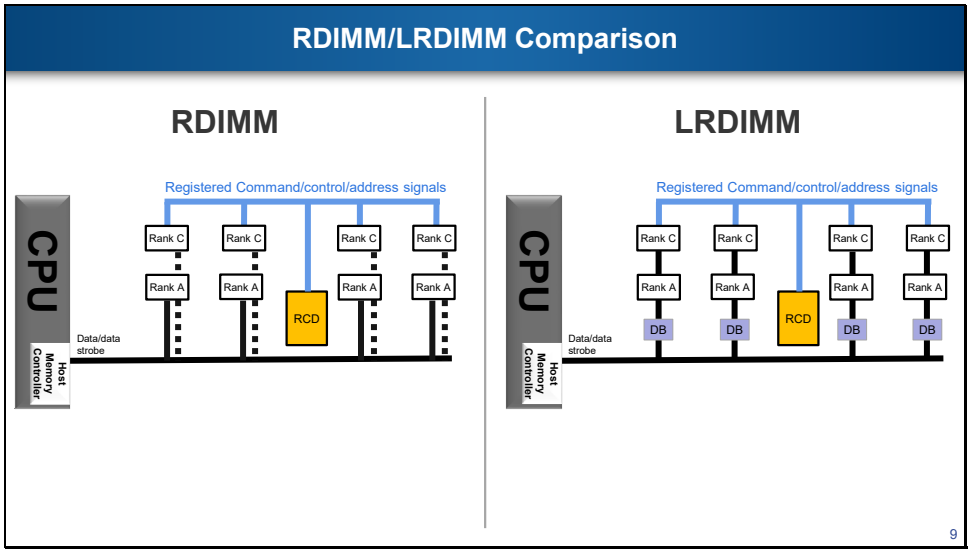

In the illustration above, the LRDIMM features multiple data buffers (DBs) to buffer incoming DQ and DQS signal lines between the host memory controller and DRAM, reducing the load on the host memory controller relative to the load experienced by the exemplary RDIMM.

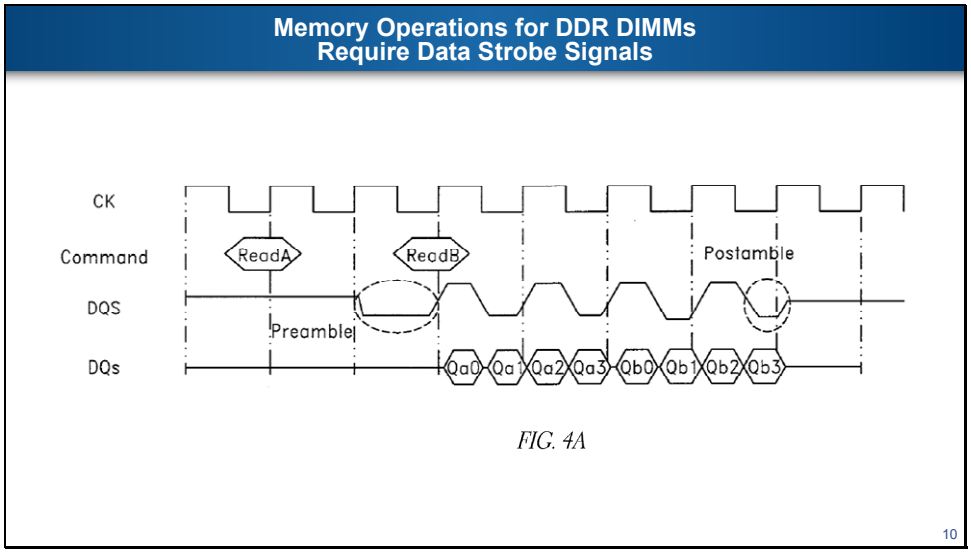

DDR stands for "double data rate." DDR DRAM uses both the rising edge and the falling edge of the clock cycle to transmit data. A drawback of the dual-edged clocking scheme of DDR is that it has fewer clock edges available to synchronize transmissions. Consequently, DDR SDRAM uses a data strobe signal to ensure accurate timing for driving and sampling data. The data strobe signal, or "DQS" signal, is specified by the DDR standard. Generally, a strobe signal indicates that another signal, e.g., data or command, is present and valid. This allows the system to distinguish each data transmission cycle accurately in a clock cycle, so as to facilitate the receiver to receive data accurately.

According to the '386 patent, "Due to their source synchronous nature, DDR SDRAM (e.g., DDR1, DDR2, DDR3) memory devices operate with a data transfer protocol which surrounds each burst of data strobes with a pre-amble time interval and a post-amble time interval. The pre-amble time interval provides a timing window for the receiving memory device to enable its data capture circuitry when a known valid level is present on the strobe signal to avoid false triggers of the memory device's capture circuit. The post-amble time interval provides extra time after the last strobe for this data capture to facilitate good signal integrity." '386 patent, 23:45-55. The above diagram shows the data strobe and data signals on host data lines. If a data buffer is placed between a host and a memory device, there will be a data path through the data buffer and the data buffer will similarly need data strobe signals in order to sample the data that passes through the data buffer.

Technology Tutorial Topics

- Overview
- **DDR4 LRDIMM**
- DDR5's on-Module Power Management
- Vertically stacked high bandwidth memory

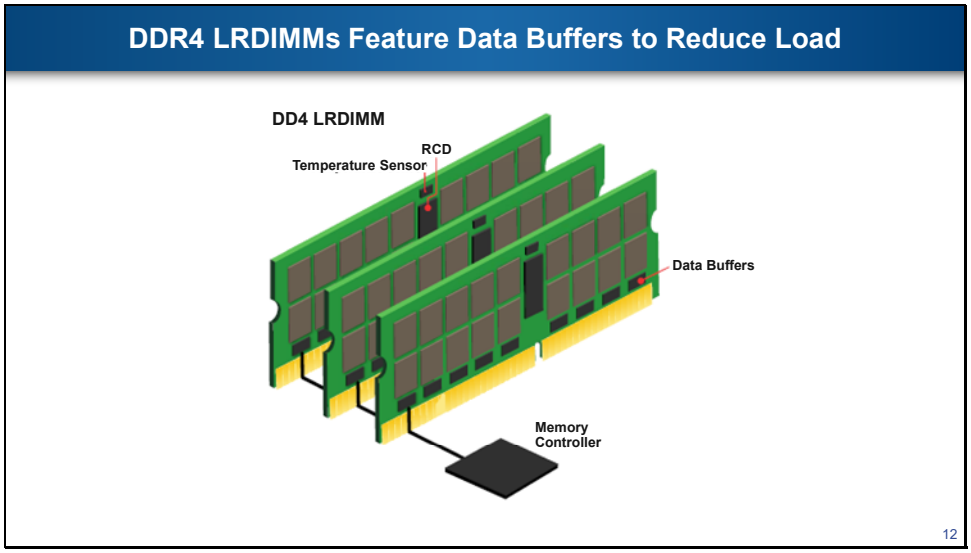

Image from: *White Paper—DDR4 LRDIMMs Let You Have It All: LRDIMMs Provide a Superior Alternative Solution for Both Deeper Memory and Higher Data Bandwidth,* https://www.eeworldonline.com/ddr4-lrdimms-let (last accessed Aug. 30, 2022)

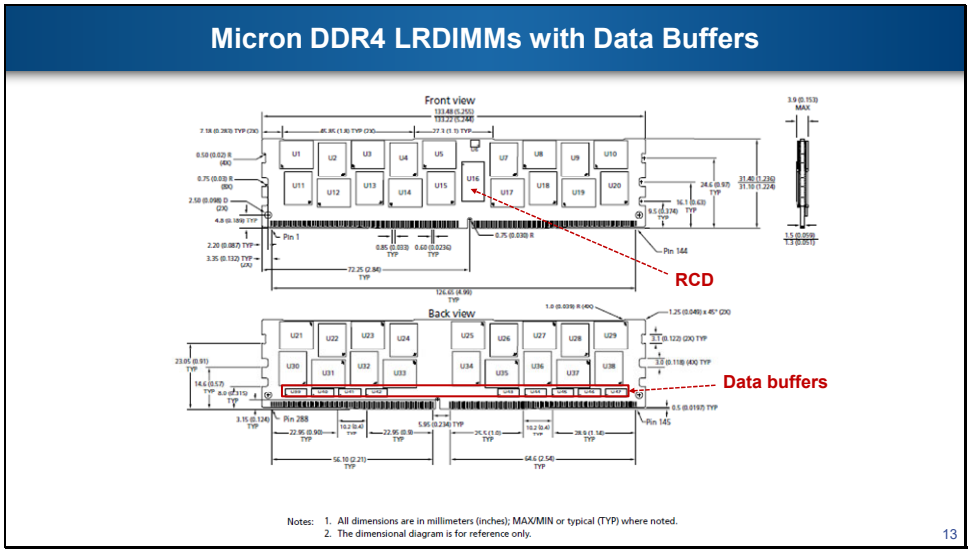

32GB (x72, ECC, QR) 288-pin DDR4 LRDIMM, MTA72ASS4G72LG-32GB, available from https://media-www.micron.com/-/media/client/global/documents/products/data-sheet/modules/lrdimm/ddr4/ass72c4gx72lz.pdf?rev=35600a94204c406f9a806cb9815055c0.

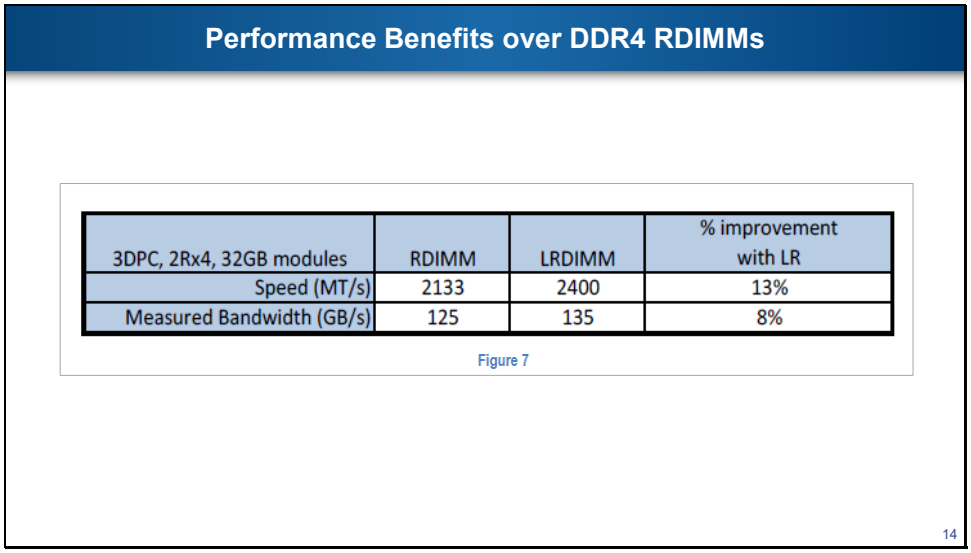

Performance Benefits over DDR4 RDIMMs

3DPC, 2Rx4, 32GB modules	RDIMM	LRDIMM	% improvement with LR
Speed (MT/s)	2133	2400	13%
Measured Bandwidth (GB/s)	125	135	8%

Figure 7

14

LRDIMMs feature higher speeds and greater bandwidth than comparable RDIMMs. *White Paper—DDR4 LRDIMMs Let You Have It All: LRDIMMs Provide a Superior Alternative Solution for Both Deeper Memory and Higher Data Bandwidth,* https://www.eeworldonline.com/ddr4-lrdimms-let (last accessed Aug. 30, 2022)

Technology Tutorial Topics

- Overview
- DDR4 LRDIMM
- **DDR5's on-Module Power Management**
- Vertically stacked high bandwidth memory

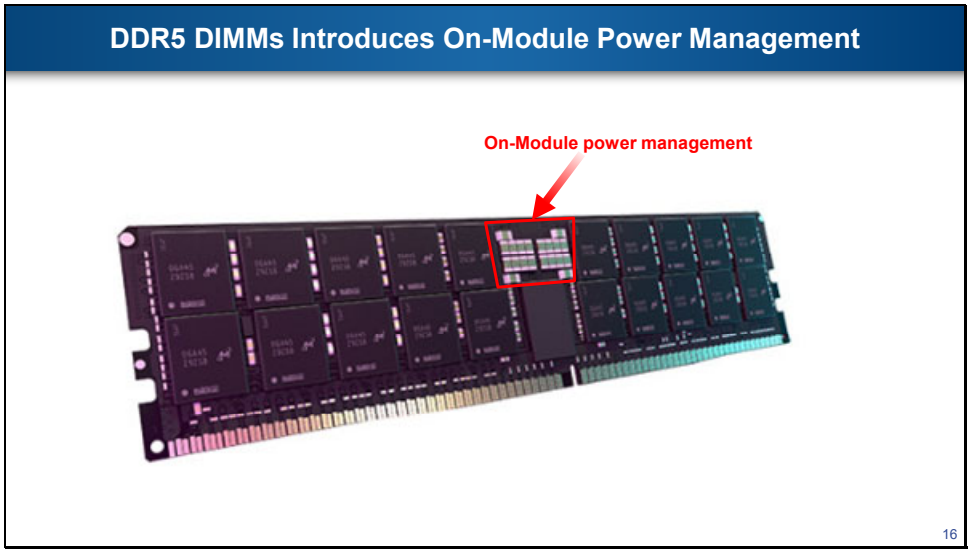

Shown above is a DDR5 RDIMM image shown at https://www.micron.com/products/dram/ddr5-sdram. The components in the red box provide on-module power management to the DIMM.

On-Module PMIC Enables Better Power Regulation

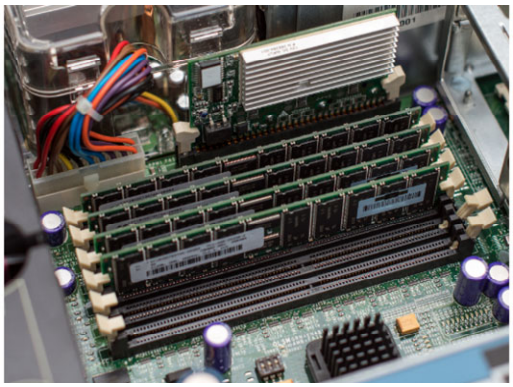

- Hundreds of wires connecting DIMMs to host's power supply to DIMMs
- Host must account for all voltage drops between host's power supply and DIMM
- As DRAM's operating voltage drops, the same 0.1V difference becomes more significant

DDR Gen.	Vdd (V)	% of error with 0.1V margin
DDR1	2.5V	4.0%
DDR2	1.8V	5.6%
DDR3	1.5V	6.7%
DDR4	1.2V	8.3%
DDR5	1.1V	9.1%

17

Image from: https://whatsabyte.com/dimm-slots.

Previous generations of memory modules (DDR4, DDR3, DDR2 …) put the power management on the motherboard, external to the memory module. There are hundreds of wires connecting the motherboard to each DIMM. When the host provides power to the memory module, various voltages are sent on selected wires/pins to provide power to corresponding components. The voltage at the CPU has to account for any voltage drop that occurs between the CPU and the memory module.

In DDR5, power management is moved onto the DIMM module board itself. The system provides a regulated voltage to the module, and the PMIC converts the provided voltage to the various regulated voltages needed to power the components of the module. As a result, the wire traces between the DIMM components and their respective power sources (i.e., the PMICs' outputs) are much shorter with on-module power management. This in turn allows for far more precise control of the voltages, given the much shorter lengths of the wire traces.

On-Module PMIC Improves Signal Qualities

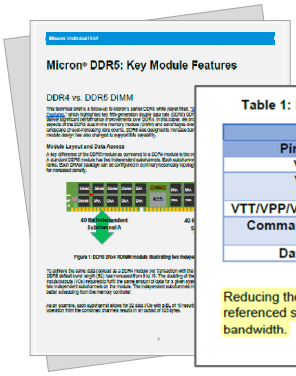
Micron® DDR5: Key Module Features

DDR4 vs. DDR5 DIMM

Table 1: Pin Differences Between DDR4 and DDR5 L/RDIMMs

Overview of L/RDIMM Pin Condition Changes			
Pin Type	DDR4	DDR5	Note
VDD	26 – VDD	3 – 12V (bulk)	Reduces overall power pins
VSS	94	127	Increased ground for SI; signals are all VSS referenced
VTT/VPP/VREF/VDDSPD	9	0	PMIC supports all these rails
Command/Address	27 + CS	2 x 7 (DDR) + CS	Seven DDR CA pins per subchannel; plus, chip select (CS) pin
Data I/Os	72	80	Support for separate subchannels

Reducing the number of power rail and CA pins allows adding more VSS ground pins to improve DDR5 VSS-referenced signal crosstalk as well as other signal-integrity challenges caused by increasing the overall system bandwidth.

18

Micron DDR5: Key Module Features, Technical Brief (2020), *available at* https://media-www.micron.com/-/media/client/global/documents/products/technical-marketing-brief/ddr5_key_module_features_tech_brief.pdf?la=en&rev=f3ca96bed7d9427ba72b4c192dfacb56.

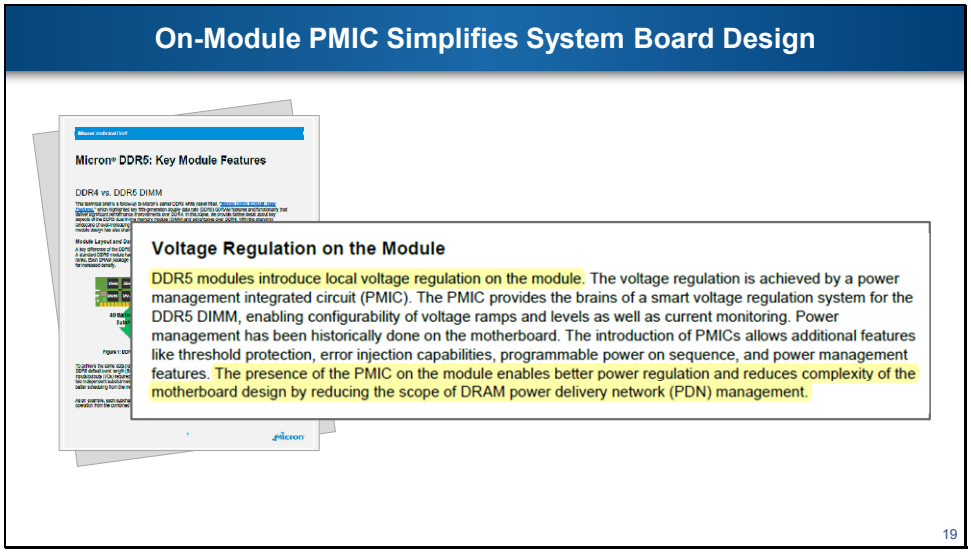

Micron DDR5: Key Module Features, Technical Brief (2020), *available at* https://media-www.micron.com/-/media/client/global/documents/products/technical-marketing-brief/ddr5_key_module_features_tech_brief.pdf?la=en&rev=f3ca96bed7d9427ba72b4c192dfacb56.

Technology Tutorial Topics

- Overview
- DDR4 LRDIMM
- DDR5's on-Module Power Management
- **Vertically stacked high bandwidth memory**

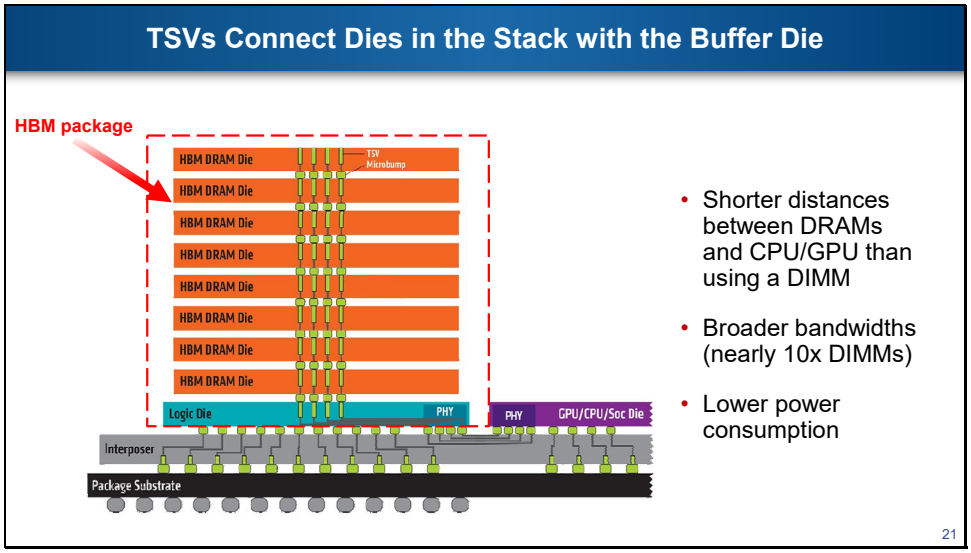

High bandwidth memory ("HBM") is a type of high-speed computer memory technology that features vertically stacked memory dies (labelled the "HBM DRAM Die" above), with a "logic die," also known as a "base die" or "control die" connected to the stacked memory dies with through silicon vias ("TSVs"). In end products incorporating HBMs, a memory host such as a CPU and/or GPU is interconnected to the logic/driver/buffer die at the bottom of each memory die stack. The combined system of CPU/GPU and the memory stack is then mounted on a substrate for use, as illustrated above on the right. This format substantially shortens the signal length between the host and the memory, which enables shorter communication time, smaller format, higher performance and lower power consumption. HBM allows for broader bandwidth (nearly 10x over DIMMs) at much lower power consumption per bit.

See Pivotal Memory Technologies Enabling New Generation of AI Workloads, Samsung, *available at* https://developer.download.nvidia.com/video/gputechconf/gtc/2019/presentation/s91028-2-pivotal-memory-technologies-enabling-new-generation-of-ai-workloads-presented-by-samsung.pdf (last accessed August 30, 2022); *What Are HBM, HBM2 and HBM2E? A Basic Definition*, Toms Hardware, *available at* https://www.tomshardware.com/reviews/glossary-hbm-hbm2-high-bandwidth-memory-definition,5889.html.

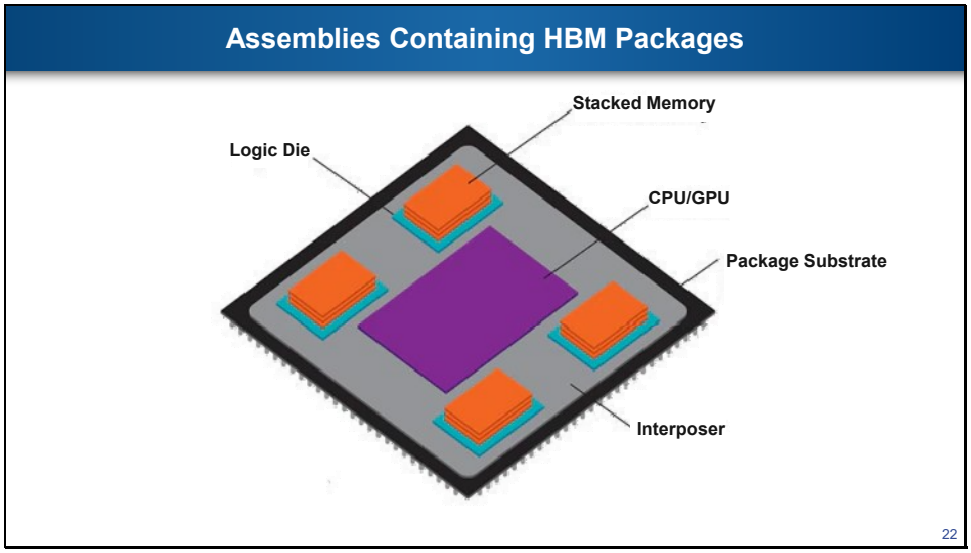

Above is an illustration of four HBM packages assembled with a host processor on a silicon interposer which is in turn mounted on a package substrate.

Source: *What Are HBM, HBM2 and HBM2E? A Basic Definition*, Toms Hardware, *available at* https://www.tomshardware.com/reviews/glossary-hbm-hbm2-high-bandwidth-memory-definition,5889.html (last accessed August 30, 2022).

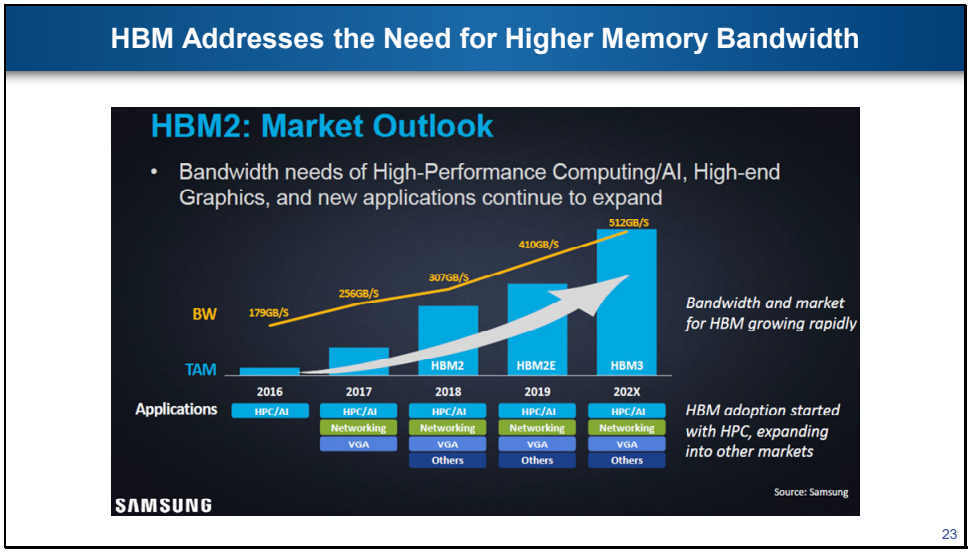

Source: *Pivotal Memory Technologies Enabling New Generation of AI Workloads*, Samsung, *available at* https://developer.download.nvidia.com/video/gputechconf/gtc/2019/presentation/s91028-2-pivotal-memory-technologies-enabling-new-generation-of-ai-workloads-presented-by-samsung.pdf (last accessed August 30, 2022).

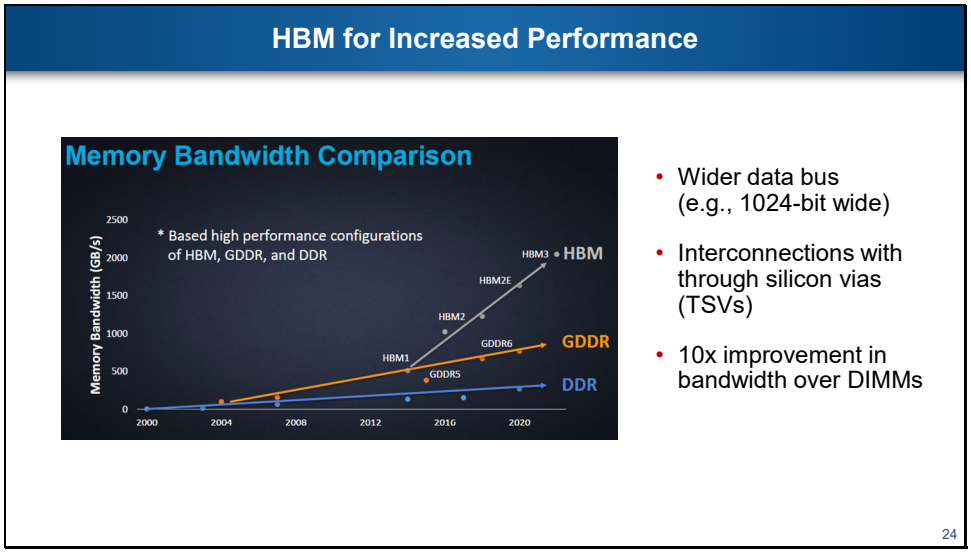

Source: *Pivotal Memory Technologies Enabling New Generation of AI Workloads*, Samsung, *available at* https://developer.download.nvidia.com/video/gputechconf/gtc/2019/presentation/s91028-2-pivotal-memory-technologies-enabling-new-generation-of-ai-workloads-presented-by-samsung.pdf (last accessed August 30, 2022).

Patents-In-Suit

Memory Modules with Distributed Data Buffer Architecture

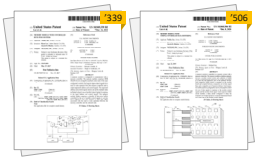

U.S. Patent **10,949,339**

U.S. Patent **10,860,506**

On-Module Power Management Integrated Circuit "PMIC" Patents

U.S. Patent 11,016,918

U.S. Patent 11,232,054

High Bandwidth Memory "HBM" Patents

U.S. Patent 8,787,060

U.S. Patent 9,318,160

The ’339 Patent

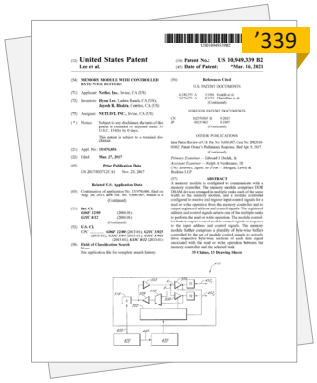

(12) **United States Patent**
Lee et al.

(10) **Patent No.: US 10,949,339 B2**
(45) **Date of Patent: *Mar. 16, 2021**

(54) **MEMORY MODULE WITH CONTROLLED BYTE-WISE BUFFERS**

(71) **Applicant: Netlist, Inc.**, Irvine, CA (US)

(72) **Inventors: Hyun Lee**, Ladera Ranch, CA (US); **Jayesh R. Bhakta**, Cerritos, CA (US)

(73) **Assignee: NETLIST, INC.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **15/470,856**

(22) Filed: **Mar. 27, 2017**

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,249,253 A 2/1981 Gentili et al.
4,571,676 A 2/1986 Mantellina et al.
(Continued)

FOREIGN PATENT DOCUMENTS

CN 102576565 B 9/2015
JP 09237492 9/1997
(Continued)

OTHER PUBLICATIONS

Inter Parties Review of U.S. Pat. No. 9,606,907, Case No. IPR2018-00362, Patent Owner's Preliminary Response, filed Apr. 9, 2017.
(Continued)

Primary Examiner — Edward J Dudek, Jr.

Problem: Expanding Memory Capacity Increases System Load

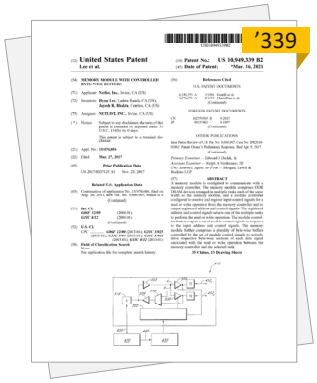

- As memory capacity increases (e.g., as the number of memory devices on a module increases) so too does the load
- Prior art arrangements with complex data line routing inefficient, limited performance

Examples of Prior Art Configurations with Load Issues

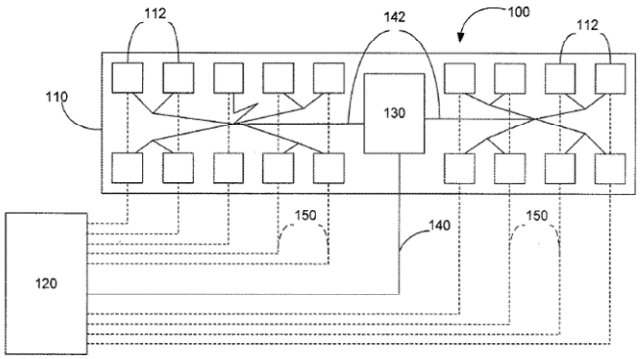

Fig. 1A

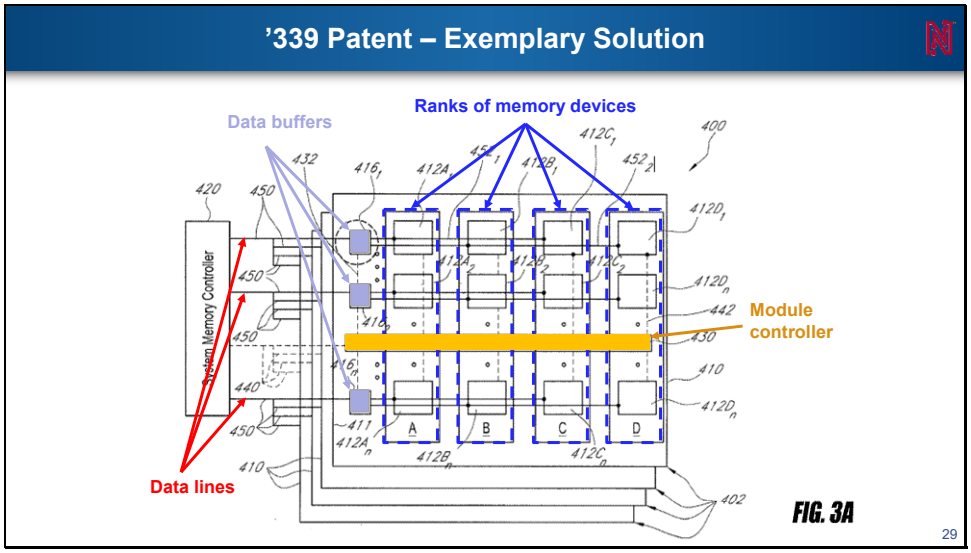

Figure 3A of the '339 patent illustrates an exemplary embodiment of the claimed memory module. The memory module includes memory devices (412) that are mounted on a printed circuit board. In the image above, the memory devices are organized in "ranks," where each "rank" is the same width as the memory module (for example, if a memory module is 64-bits wide, each rank on the memory module is 64-bits wide). The module also includes a module controller 430 configured to receive input address and control signals (440) for a read or write operation from the host's memory controller and to output registered address and control signals to the SDRAMs. As shown above, data buffers (416) are coupled between the memory controller and memory devices in each of the ranks on the memory module via data lines (450). Though Figure 3A depicts a memory module featuring four ranks of memory A-C, the '339 patent is not limited to four-rank memory modules. *See* '339 patent, 7:44-8:53; 9:44-49.

Solution Can Also Benefit Two-Rank Embodiments

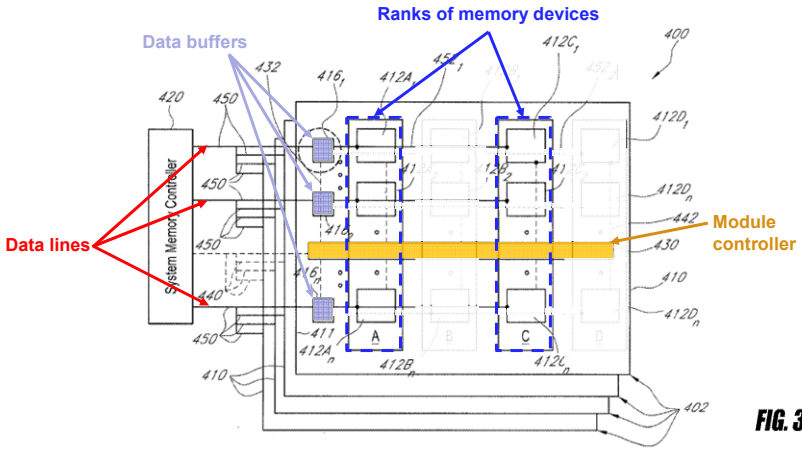

FIG. 3A

Operation at Data Buffers

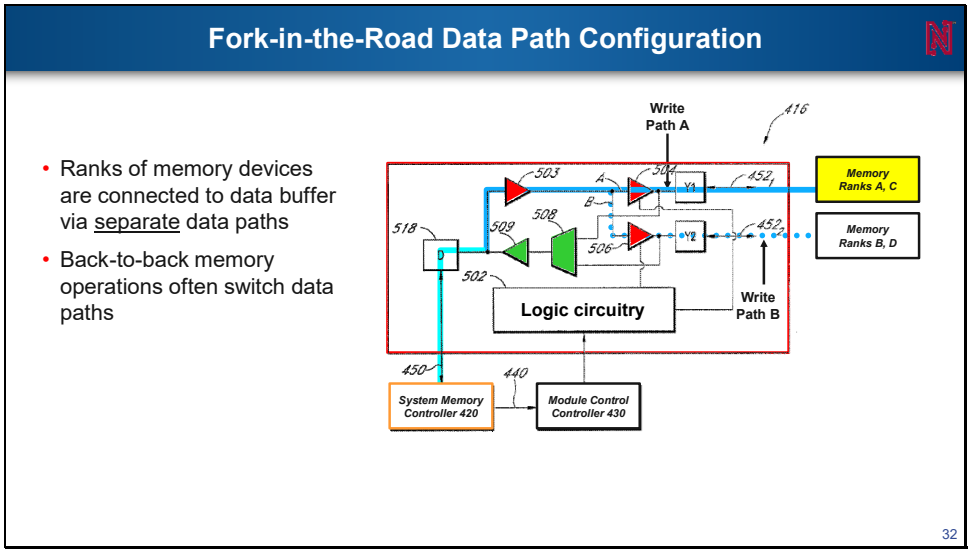

In the "Fork-in-the Road" arrangement depicted above, memory ranks A and C are connected to the data lines 452_1 on write path A while memory ranks B and D are connected to the data lines 452_2 on write path B. '339 patent, Fig. 5 (annotated).

In a write operation, in response to a signal from the module control circuit 430, the logic circuitry 502 selects either path A or path B. For example, in response to an "enable A" signal, a first tristate buffer 504 in path A is enabled to actively drive data on path A to memory rank A or C at the Y1 terminal. During this period, path B and buffer 506 are disabled. Similarly, if "enable B" is received, then tristate buffer 506 in path B is enabled to drive data along path B to memory rank B or D connected to the Y2 terminal. '339, patent, 16:7-25.

On the read path, the data passing through Y1 and Y2 are fed to a multiplexer 508 and selected by a logic circuitry 502. The selected data is then driven across buffer 509 to the system memory controller 420 across data line 518. '339, patent 16:26-44.

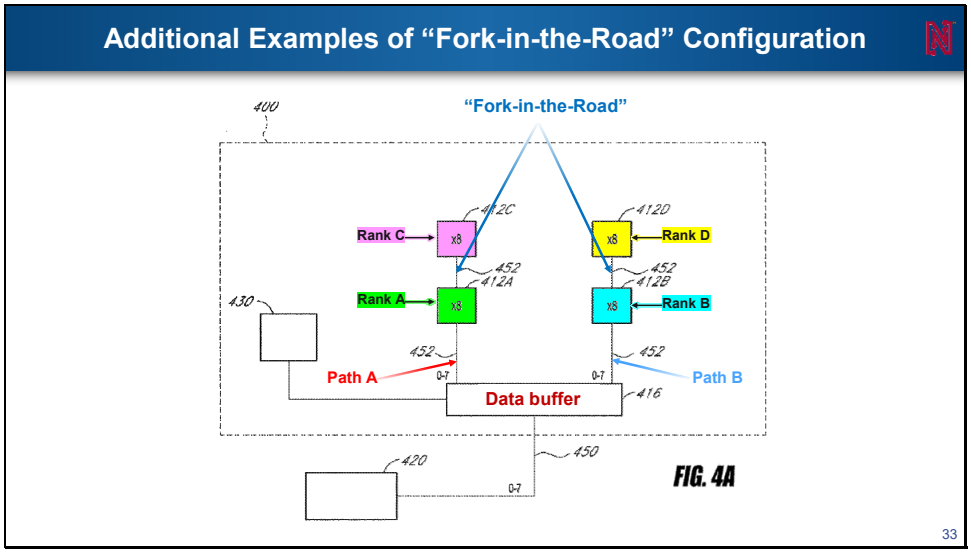

In the "Fork-in-the Road" arrangement depicted above, ranks A and C share a first data path (452) that is separate from the second data path shared by Ranks B and D. In this way, the data paths are arranged in a "Fork-in-the-Road." '339 Patent, 11:4-11:44.

Additional Examples of "Fork-in-the-Road" Configuration

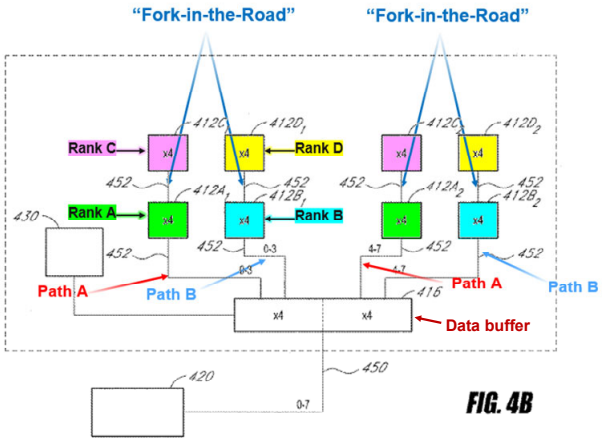

FIG. 4B

Two-Rank Embodiments May Adopt "Straight Line" Layout

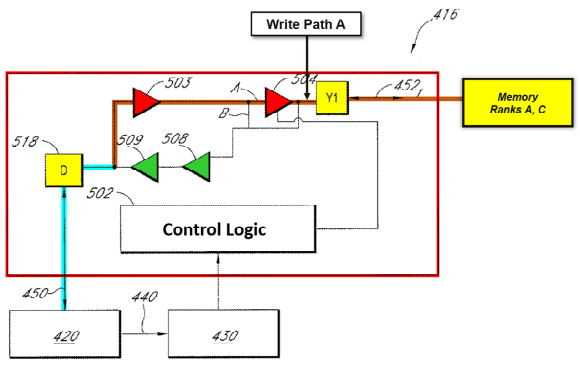

Configuration for a Two-Rank DIMM

'339 Patent, Fig. 5

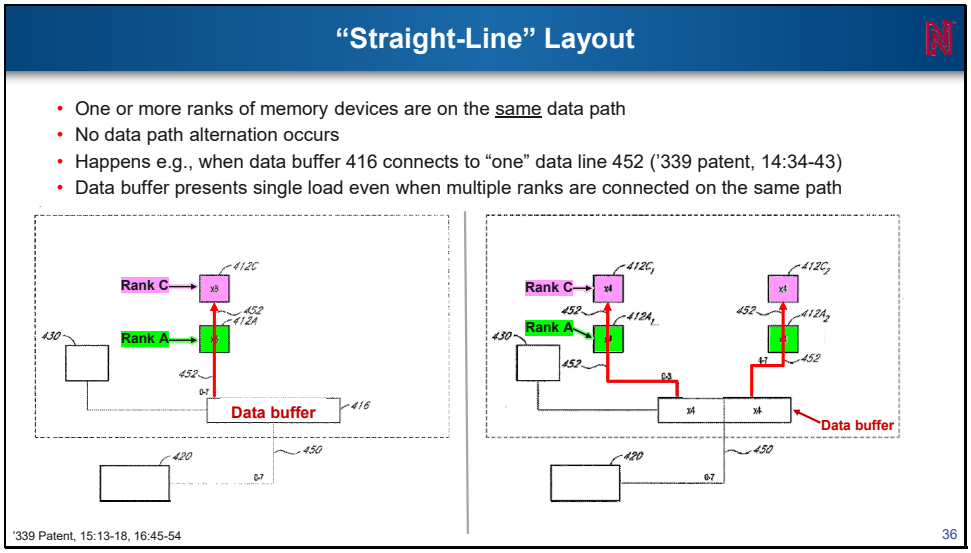

In the "Straight Line" arrangement, ranks of memory devices are on the same data path without any "Fork in the Road." This is depicted in the modified version of Figures 4A and 4B in the '339 patent. In modified Figure 4A, there is a single path between the x8 buffer and the x8 memories. In modified figure 4B, each x4 section of data is transmitted to corresponding memory devices in Rank A and Rank C. *See* '339 patent, 14:34-43.

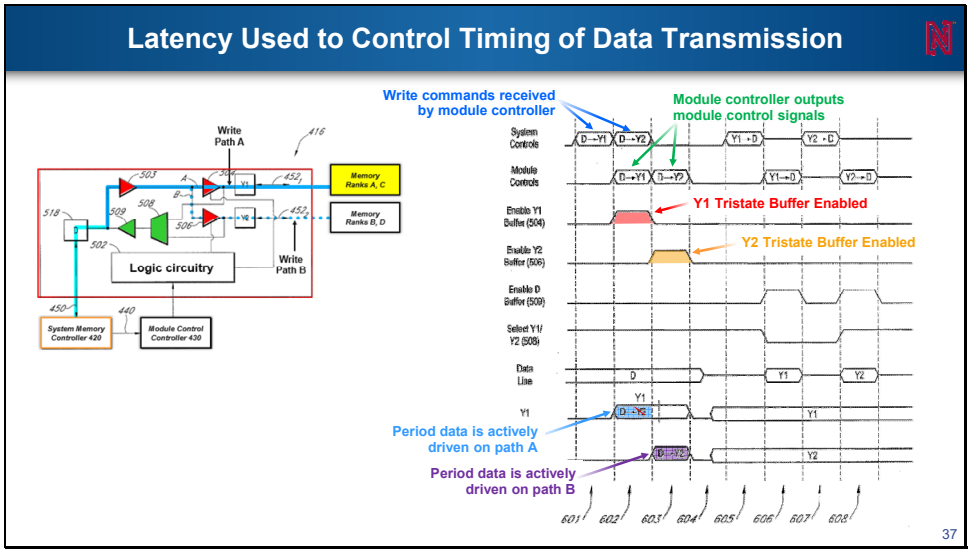

Pictured above (right) is Fig. 6 of the ’339 patent, annotated to illustrate the transmission of signals and data for a write operation during time periods 601–603. *See* ’339 patent, 17:66-18:24.

NB: Fig. 6 as depicted above is annotated to correct an obvious error: “D→Y2” on line Y1 should read “D →Y1,” consistent with the specification. *See* ’339 patent, 17:66-18:17.

Operation Related to Module Controller

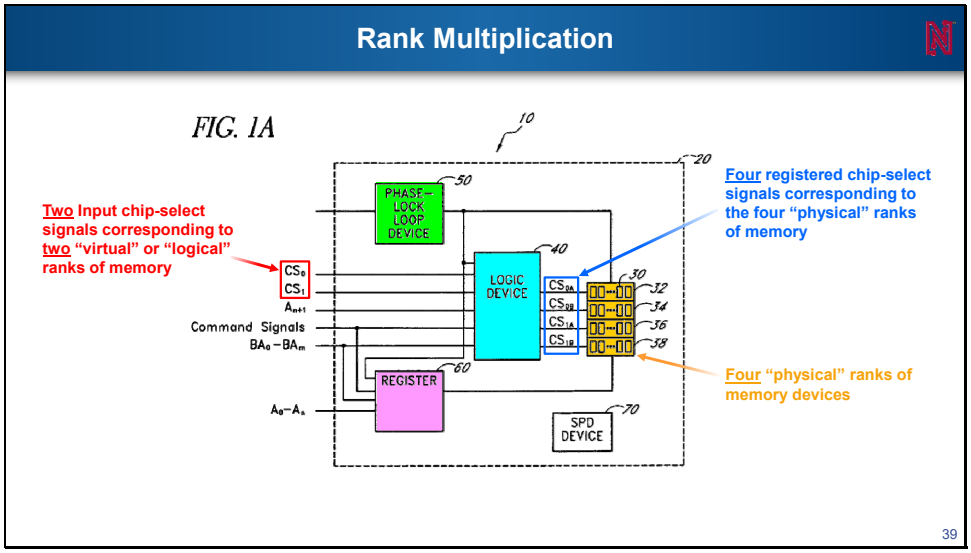

Rank multiplication is one way to generate additional chip-select signals: As depicted above, the logic element (blue) receives a set of input signals including two chip-select signals (CS_0, and CS_1), each corresponding to a single rank. The number of input chip-select signals corresponds to a number "logical" or "virtual" ranks (e.g., 2), so called because, in reality, the memory module actually has four "physical" ranks (highlighted in orange). The four "physical" ranks depicted in Figure 1A generally correspond to four chip-select signals generated by the logic element: CS_{0A}, CS_{0B}, CS_{1A}, CS_{1B}. In this way, the memory module may simulate a virtual memory module when the number of memory devices 30 of the memory module 10 is larger than the number of memory devices 30 per memory module for which the computer system is configured to utilize. In other words, from the computer system's perspective, it is connected to only two ranks of memory devices, to be selected by CS_0 or CS_1, even though the memory devices are arranged in four physical ranks. In this way, the physical (actual) number of ranks is said to be "transparent" to the computer system. Put it another way, rank multiplication allows multiple "physical" ranks of memory devices be presented as, for example, a single "logical" rank to the host system.

See '386 patent (incorporated by '339), 6:63-7:62, Fig. 1A.

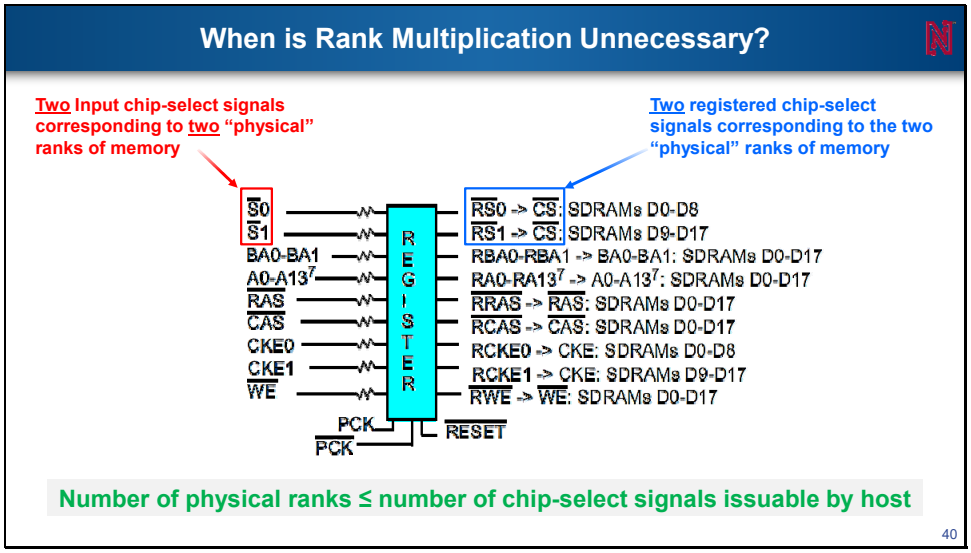

The module controller of the '339 patent may "register[] signals from the control lines 440, 440' in a manner functionally comparable to the address register of a conventional RDIMM." '339 patent, 10:38-41. The register of a conventional RDIMM, shown here as depicted in the JEDEC DDR SDRAM Registered DIMM Specification (Jan. 2002), receives a set of input address and control signals from the memory controller of the host system, which includes two chip-select signals S0 and S1. The register outputs a set of registered address and control signals, which incudes registered chip-select signals RS0 and RS1. As explained in the previous slide, rank multiplication allows a memory controller configured for a first number of chip-select signals (e.g., 2) corresponding to a "virtual" or "logical" number of ranks, to emulate a system that is configured for a greater number of chip-select signals (e.g., 4) corresponding to a "physical" number of ranks (i.e., the number of ranks the module actually has). In a conventional memory module without rank multiplication, the number of input chip-select signals and the number of registered chip-select signals both correspond to the number of "physical" ranks (which is 2 in the above illustration).

’506 Patent

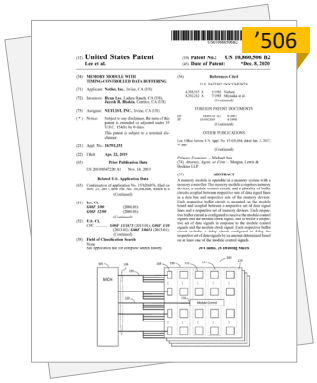

(12) **United States Patent**
Lee et al.

(10) **Patent No.: US 10,860,506 B2**
(45) **Date of Patent: *Dec. 8, 2020**

(54) **MEMORY MODULE WITH TIMING-CONTROLLED DATA BUFFERING**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

(72) Inventors: **Hyun Lee**, Ladera Ranch, CA (US); **Jayesh R. Bhakta**, Cerritos, CA (US)

(73) Assignee: **NETLIST, INC.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **16/391,151**

(22) Filed: **Apr. 22, 2019**

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,368,515 A 1/1983 Nielsen
4,392,212 A 7/1983 Miyasaka et al.
(Continued)

FOREIGN PATENT DOCUMENTS

EP 1816570 A2 8/2007
JP 10-092160 4/1998
(Continued)

OTHER PUBLICATIONS

Lee, Office Action, U.S. Appl. No. 15/426,064, dated Jun. 2, 2017, 11 pgs.
(Continued)

Problem: Read/Write Leveling Insufficient to Ensure Strict Timing Constraints Are Met

- Background: Proper timing of data transmission requires that the distribution of control/clock signals in the memory module be subject to strict constraints.
- Conventional solution: the memory controller performs read/write leveling techniques to time transmissions of data.
- But in a distributed data buffer architecture, data buffers are in the data path between the memory controller and DRAMs → read/write leveling insufficient.

42

'506 patent, 2:16-36, 15:17-26.

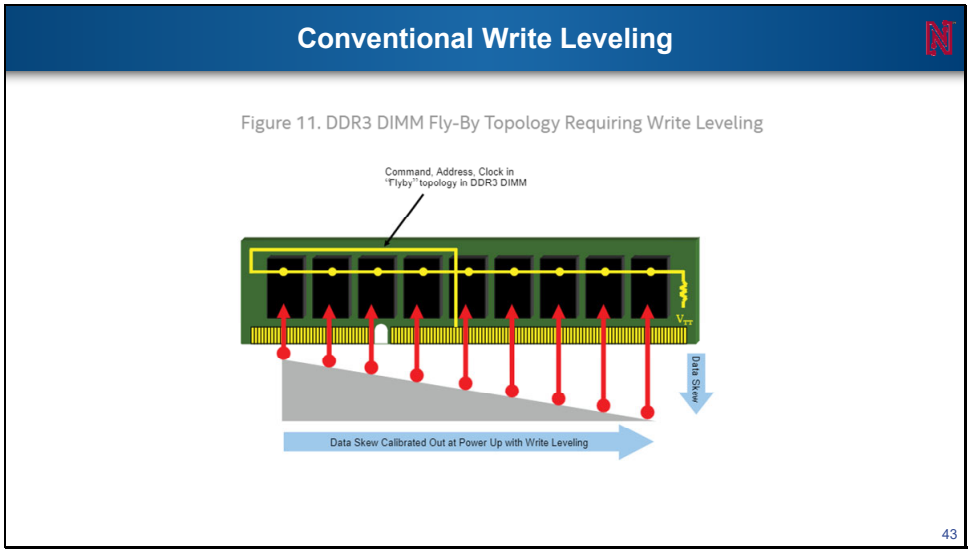

Write leveling was introduced in DDR3 as one solution to flight-time skew in the "fly-by" memory module topology depicted above. *See* External Memory Interface Handbook Volume 2: Design Guidelines: For UniPHY-based Device Families, Intel, § 2.1.1, available at https://www.intel.com/content/www/us/en/docs/programmable/683385/17-0/read-and-write-leveling.html.

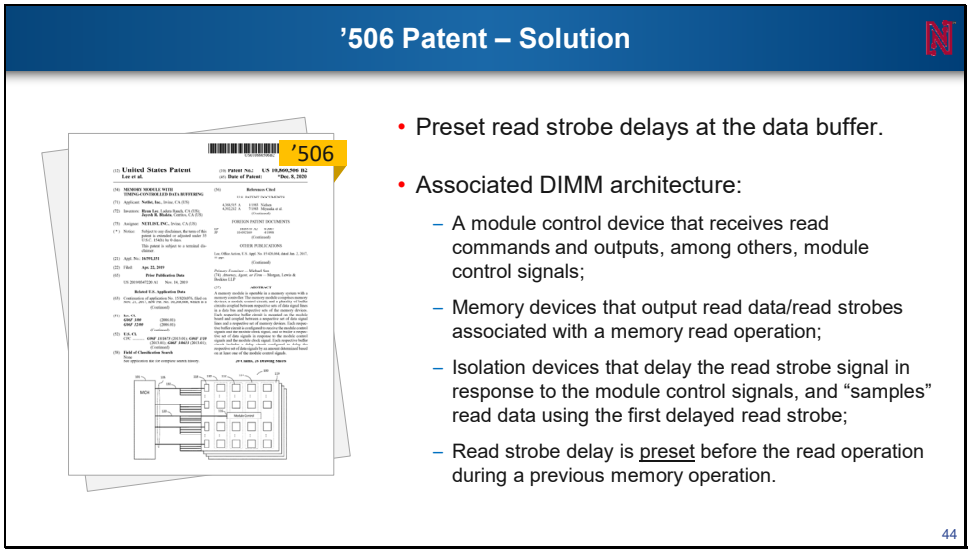

The '506 patent solves the problem of timing data transmission for a read command in the prior art by shifting the burden of managing timing of data transmission during a read command away from the memory controller and onto the data buffers. The claims of the '506 patent are directed to a memory module featuring memory devices, a module controller, and data buffers that include circuitry configured to delay data associated with a read command by reference to a delayed read strobe signal. Generally, a "strobe" signal is a signal that indicates that another signal, e.g., data or command, is present and valid, and can be used to accurately time the transmission of data throughout the module. Bruce Jacob et al., Memory Systems: Cache, DRAM, Disk, 318 (2008). The delayed read strobe signal is delayed by a predetermined amount determined based on at least one of the module control signals received by the data buffer. *See, e.g.*, '506 cl. 1.

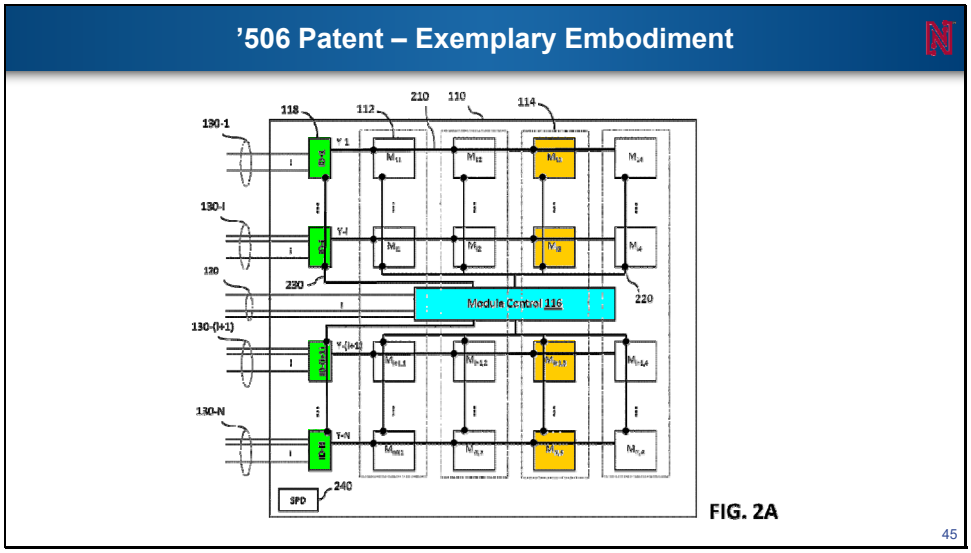

As shown in FIG. 2A, the module control device (116) receives control/address signals (120) from a the computer system's memory controller (not shown). The module control device is coupled to each of the memory devices arranged in multiple ranks (114), and data buffers (118). In operation, the module control device receives input C/A signals corresponding to a memory operation via the C/A signal line, and transmits registered C/A signals and module control signals to the memory devices 112 and the buffers 118, respectively. *See* '506, 5:11-26, 60-62, 7:66-8:3.

Patents-In-Suit

Memory Modules with Distributed Data Buffer Architecture

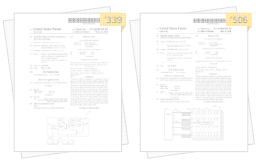

U.S. Patent 10,949,339

U.S. Patent 10,860,506

On-Module Power Management Integrated Circuit "PMIC" Patents

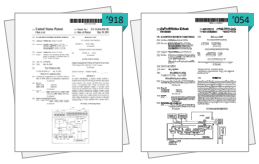

U.S. Patent **11,016,918**

U.S. Patent **11,232,054**

High Bandwidth Memory "HBM" Patents

U.S. Patent 8,787,060

U.S. Patent 9,318,160

’918 Patent

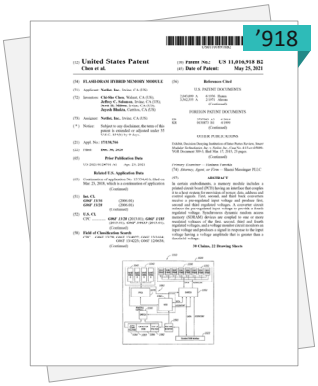

(12) **United States Patent**
Chen et al.

(10) **Patent No.: US 11,016,918 B2**
(45) **Date of Patent: May 25, 2021**

(54) **FLASH-DRAM HYBRID MEMORY MODULE**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

(72) Inventors: **Chi-She Chen**, Walnut, CA (US); **Jeffrey C. Solomon**, Irvine, CA (US); **Scott H. Milton**, Irvine, CA (US); **Jayesh Bhakta**, Cerritos, CA (US)

(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **17/138,766**

(22) Filed: **Dec. 30, 2020**

(56) **References Cited**

U.S. PATENT DOCUMENTS

2,043,099 A 6/1936 Hanna
3,562,555 A 2/1971 Ahrons
(Continued)

FOREIGN PATENT DOCUMENTS

EP 2737383 A2 6/2014
KR 0130873 B1 4/1999
(Continued)

OTHER PUBLICATIONS

Exhibit, Decision Denying Institution of Inter Partes Review, *Smart Modular Technologies, Inc.* v. *Netlist, Inc.*, Case No. 4:13-cv-05889-YGR Document 309-3, filed Mar. 17, 2015, 23 pages.

’054 Patent

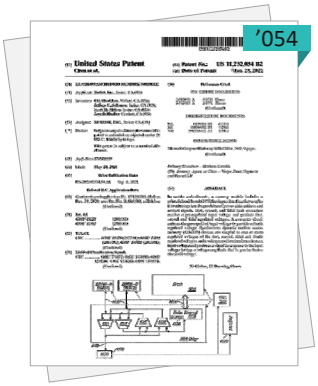

(12) **United States Patent**
Chen et al.

(10) **Patent No.: US 11,232,054 B2**
(45) **Date of Patent: *Jan. 25, 2022**

(54) **FLASH-DRAM HYBRID MEMORY MODULE**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)

(72) Inventors: **Chi-She Chen**, Walnut, CA (US); **Jeffrey C. Solomon**, Irvine, CA (US); **Scott H. Milton**, Irvine, CA (US); **Jayesh Bhakta**, Cerritos, CA (US)

(73) Assignee: **NETLIST, INC.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **17/328,019**

(22) Filed: **May 24, 2021**

(56) **References Cited**

U.S. PATENT DOCUMENTS

2,043,099 A 6/1936 Hanna
3,562,555 A 2/1971 Ahrons
(Continued)

FOREIGN PATENT DOCUMENTS

KR 0130873 B1 4/1999
KR 100600242 B1 7/2006
WO 2013016773 A2 1/2013

OTHER PUBLICATIONS

Microsoft Computer Dictionary Fifth Edition, 2002, 9 pages.
(Continued)

Primary Examiner — Hashem Farrokh

Problem – Power Supply to Volatile Memory Systems

- Power regulation is critical to operations: Volatile memory (e.g., SDRAMs on DIMMs) maintains stored information only when it is powered
- Batteries are a conventional alternative to powering volatile memory in the event of a power fluctuation, but require maintenance, are not environmentally friendly, and the status of batteries can be difficult to monitor

49

See '918 patent, 3:53-62, 24:9-16, 24:60-25:8, 26:4-8.

PMIC Patents – Solutions

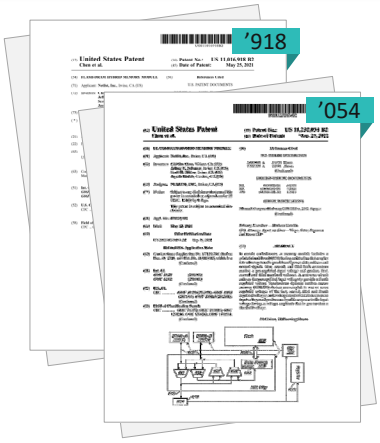

- On-DIMM power management integrated circuit that provides a plurality of regulated voltages to the various components on the DIMM (e.g., SDRAMs, data buffers, RCD, etc.);
- On-DIMM voltage monitoring circuit:
  - Monitors an input voltage and produces a trigger signal in response to the input voltage having a voltage amplitude that is less than or greater than a threshold voltage
  - Causes information to be written to the non-volatile memory, including in controller

50

See, e.g., '918, Abstract, 20:43-57, 26:54-26:3, 29:18-64, Fig. 16.

Exemplary Embodiment

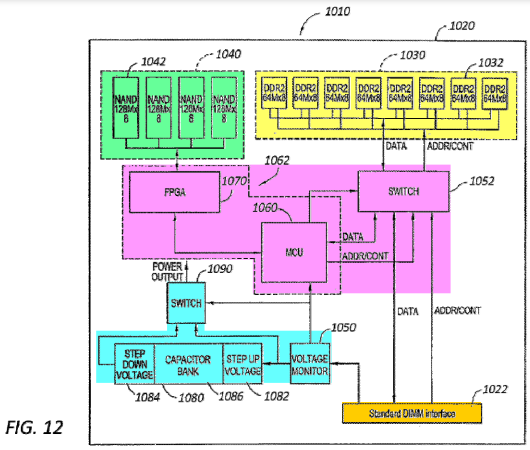

FIG. 12

51

In the above example, a first power supply is provided via the DIMM interface 1022 (orange). There is also a second power supply 1080 (colored in blue) for use when a trigger event occurs (*e.g.*, system hang-up or power failure). The second power supply is connected to the rest of the system via a switch 1090. There can also be a third power supply (not shown) for use when it is detected that a trigger event is likely to occur but has not yet occurred.

The rest of the DIMM system includes memory (volatile and/or non-volatile), controller 1062 which can include both a microcontroller 1060 and a logic element (FPGA) 1070. Switch 1052 operatively couples or decouples volatile memory from the controller 1062.

See '918 patent, 23:1-40, 25:54- 26:3, 26:26-35, 26:36-65.

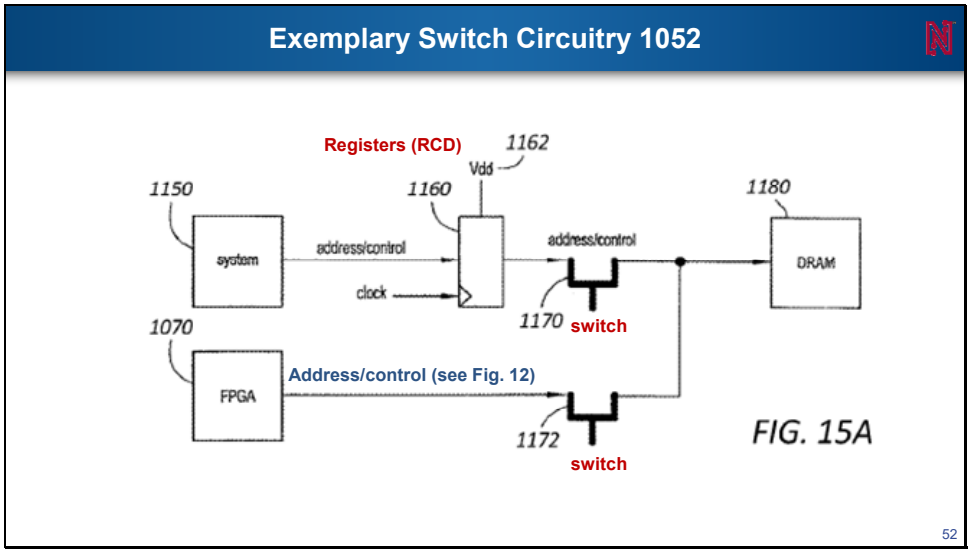

'918, 23:41-56. The above shows an example of a switch circuitry 1052 for a memory system comprising a registered DIMM ("RDIMM") subsystem. 1160 is one or more registers, which would be understood as the RCD of the RDIMM. RCD is understood as a controller for a registered DIMM (RDIMM). The switch circuitry also includes switches 1170 and 1172. 1160 is powered by Vdd, which could be 1.8V for a DDR2 system. The switches can be supplied with a voltage at a different voltage level, such as 2.5V, 3.3V or another voltage. 29:46-54.

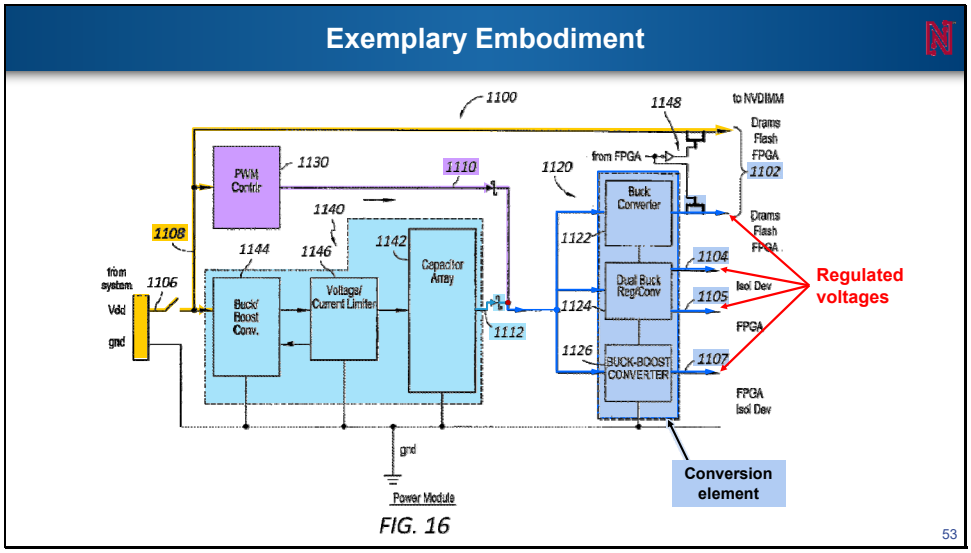

Figure 16 provides a more detailed implementation of the power supplies shown in Figure 12. In particular, the depicted power module 1100 includes a second power element 1140 that corresponds to the second power supply 1080 in Figure 12, for use when a trigger event has occurred. A first power element 1130 corresponds to the third power supply referenced for Figure 12, for use when the memory system detects that a trigger signal is likely to occur but has not occurred. A system power supply, corresponding to Figure 12's first power supply, provides an input voltage 1106 that is transformed to voltage 1108. Voltage 1108 supplies power to the power module 1100. Voltage 1108 is fed into both the first and second power elements to generate regulated voltages 1110 and 1112 for input to power conversion circuit 1120. Voltage 1108 can also directly power memory and FPGA, as voltage 1102, when no trigger event is detected or has occurred. Otherwise, memory and FPGA can be powered by voltage 1102 output by buck converter 1122. While in Figure 16, the volatile and non-volatile memories are powered by a single voltage 1102, they may also be powered by independent voltages.

Buck converter 1122 is part of the conversion element 1120. Other sub-blocks in conversion element 1120 may include (i) a dual buck regulator/converter 1124, which outputs voltages 1104 and 1105, and (ii) a buck-boost converter 1126 that outputs a voltage 1107. Voltage 1104 can be used to power an isolation device, voltage 1105 can be used to power controller 1162 and voltage 1107 can be used to power both the controller and the isolation circuit. The PMIC patents do not place a restriction on the value of the output voltages. *See* 27:59-28:55; 29:18-54, 29:61-64.

Known Types of Voltage Conversion Circuits

- Buck converter

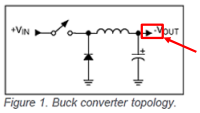

Figure 1. Buck converter topology.

- Linear dropout regulator (LDO)

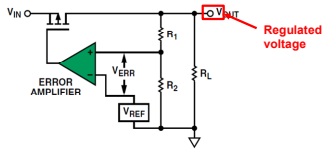

- Dual-buck converter

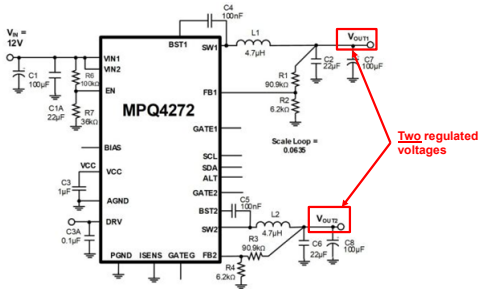

54

Sources: https://www.maximintegrated.com/en/design/technical-documents/tutorials/2/2031.html (buck converter); https://www.analog.com/en/analog-dialogue/articles/low-dropout-regulators.html (linear dropout regulator); https://www.monolithicpower.com/en/mpq4272-aec1.html (dual buck converter).

Patents-In-Suit

Memory Modules with Distributed Data Buffer Architecture

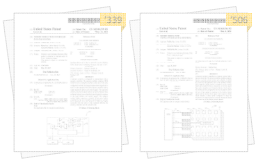

U.S. Patent 10,949,339

U.S. Patent 10,860,506

On-Module Power Management Integrated Circuit "PMIC" Patents

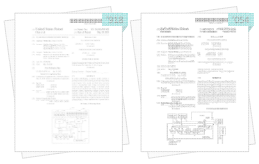

U.S. Patent 11,016,918

U.S. Patent 11,232,054

High Bandwidth Memory "HBM" Patents

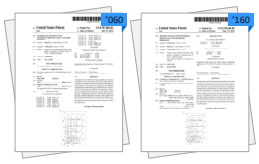

U.S. Patent **8,787,060**

U.S. Patent **9,318,160**

’060 Patent

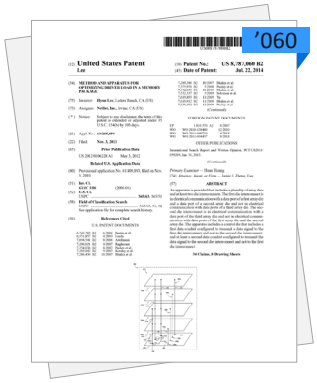

(12) **United States Patent**
Lee

(10) **Patent No.: US 8,787,060 B2**
(45) **Date of Patent: Jul. 22, 2014**

(54) **METHOD AND APPARATUS FOR OPTIMIZING DRIVER LOAD IN A MEMORY PACKAGE**

(75) Inventor: **Hyun Lee**, Ladera Ranch, CA (US)

(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 103 days.

(21) Appl. No.: **13/288,850**

(22) Filed: **Nov. 3, 2011**

7,289,386	B2	10/2007	Bhakta et al.
7,375,970	B2	5/2008	Pauley et al.
7,442,050	B1	10/2008	Bhakta et al.
7,532,537	B2	5/2009	Solomon et al.
7,619,893	B1	11/2009	Yu
7,619,912	B2	11/2009	Bhakta et al.
7,630,202	B2	12/2009	Pauley et al.

(Continued)

FOREIGN PATENT DOCUMENTS

EP	1 816 570	A2	8/2007
WO	WO 2010-138480		12/2010
WO	WO 2011-049710		4/2011
WO	WO 2011-094437		8/2011

OTHER PUBLICATIONS

’160 Patent

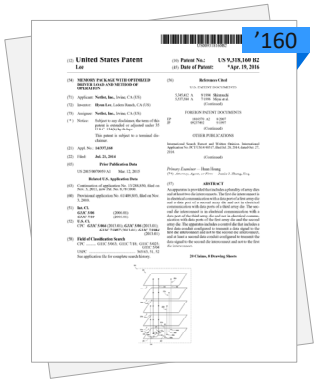

(12) **United States Patent**
Lee

(10) **Patent No.: US 9,318,160 B2**
(45) **Date of Patent: *Apr. 19, 2016**

(54) **MEMORY PACKAGE WITH OPTIMIZED DRIVER LOAD AND METHOD OF OPERATION**

(71) Applicant: **Netlist, Inc.**, Irvine, CA (US)
(72) Inventor: **Hyun Lee**, Ladera Ranch, CA (US)
(73) Assignee: **Netlist, Inc.**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **14/337,168**
(22) Filed: **Jul. 21, 2014**

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,345,412 A 9/1994 Shiratsuchi
5,537,584 A 7/1996 Miyai et al.
(Continued)

FOREIGN PATENT DOCUMENTS

EP 1816570 A2 8/2007
JP 09237492 9/1997
(Continued)

OTHER PUBLICATIONS

International Search Report and Written Opinion, International Application No. PCT/US14/48517, filed Jul. 28, 2014, dated Oct. 27, 2014.
(Continued)

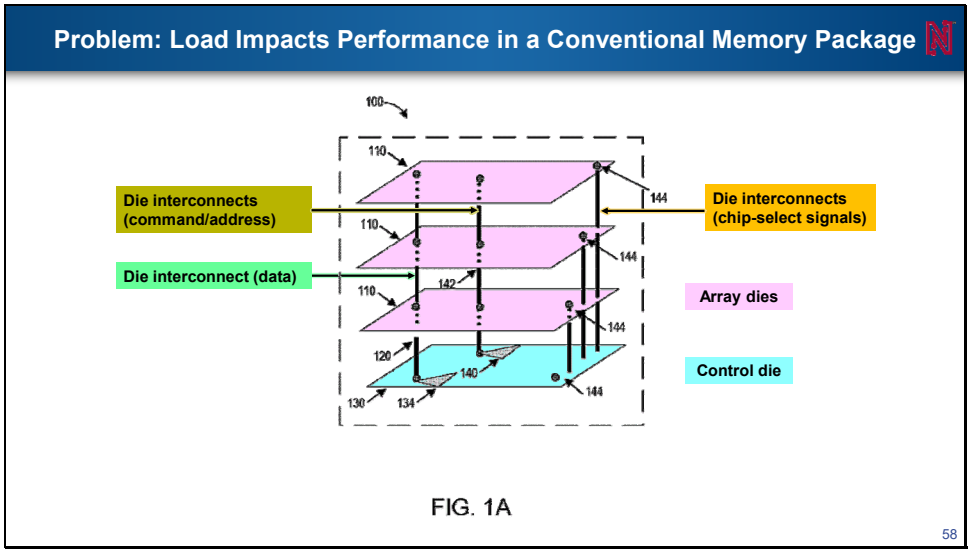

Figure 1A from the '060/'160 patents depicts then-existing memory packages comprised of a control die and a plurality of stacked array dies connected to the control die via die interconnects. In those conventional packages, the die interconnects are in communication with each of the array dies, which disadvantageously increases the load on the drivers, requiring larger driver sizes which increase the size of the control die and consumes more power. *See* '060, 1:30-56; 2:8-15.

HBM Patents (’060/’160) – Solution

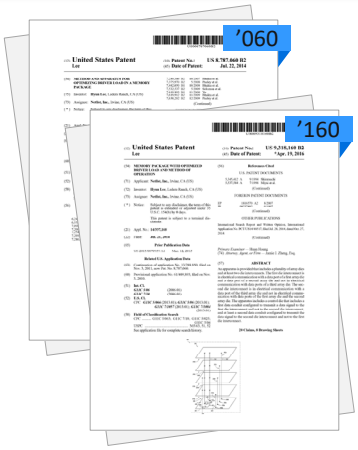

- Load Reduction: die interconnects in electrical communication with some, but not all of the array dies.
- Associated structural change in control dies:
  - Data conduits in connection with respective die interconnects
  - Control conduit to control respective states of the data conduits in response to control signals received from host
  - Differential driver sizes for different data conduits

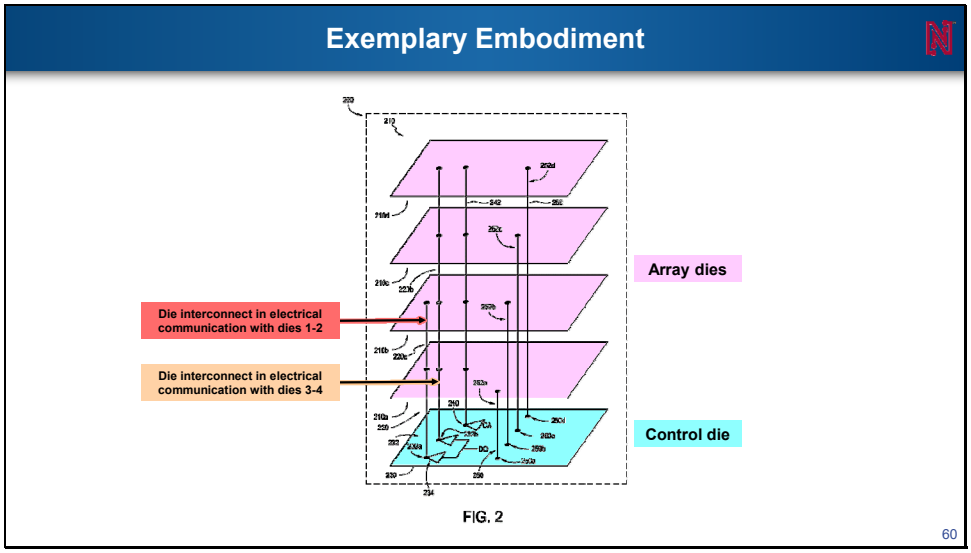

As depicted in the exemplary embodiment, Figure 2 from the '060/'160 patents, the memory package includes die interconnects in electrical communication with some but not all of the array dies. Those die interconnects that are in electrical communication with the data ports of the array dies are illustrated by the darkened circles above, while those die interconnects that are not in electrical communication with the data ports of the array dies are illustrated by the unfilled circles. As shown above, at least one of the die interconnects is in electrical communication with at least one data port from each of at least two array dies without being in electrical communication with a port from at least one array die, which may be in electrical communication with a different die interconnect. This enables the memory packages to be designed with smaller form factor in mind, and lowers power consumption. *See* '060 patent, 5:54-62, 7:22-8:62.

In the '060 patent, "data ports enable electrical communication and data transfer between the corresponding memory circuitry of the array dies 210 and a communication pathway (e.g., a die interconnect)." '060, 5:42-45. There is no electrical communication between a die interconnect and an array die when the TSV "pass[es] through" the array die such that the interconnect does "not enable electrical communication between the die interconnect … and data ports of the array die[]," as illustrated by unfilled circles. *Id.,* 8: 35-42. This may be provided by "an insulator or an air gap between the die interconnect and array die circuitry that is large enough to prevent electrical communication between the die interconnect and the array die circuitry." *Id.*, 8:47-53. There may be electrical connections between a die interconnect and an array die that the die interconnect is not configured to be in communication with: "[E]lectrical connections leading from the TSV of the array dies that are not configured to be in communication with the die interconnect may not exist or may be stubs. These stubs are not configured to provide electrical communicaton with the memory cells of the array die." *Id.*, 8:57-62.